Generating Topical Poetry

Marjan Ghazvininejad
Xing Shi, Yejin Choi, Kevin Knight

Hafez: a Poet From Shiraz, Iran

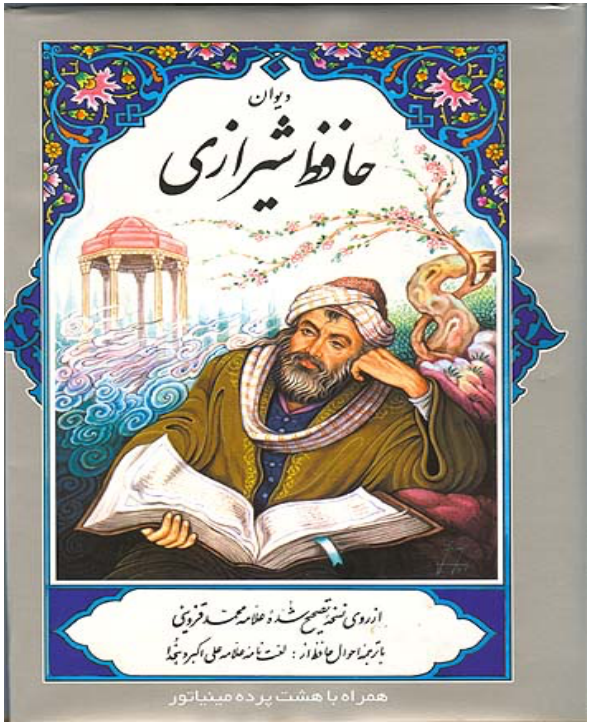

Hafez: a Poet From ISI

A system that creates computer generated poems

- Any user-supplied topic

- Any number of distinct poems on that topic

- 20-30 seconds per poem

Challenges

- Any user-supplied topic

Challenges

- Any user-supplied topic
  Like: burritos, Turing test, unicorn, Donald Trump

Challenges

- Any user-supplied topic
  Like: burritos, Turing test, unicorn, Donald Trump

- Any number of distinct poems on that topic
  One million equally good poems on french fries

Challenges

- Any user-supplied topic
  Like: burritos, Turing test, unicorn, Donald Trump

- Any number of distinct poems on that topic
  One million equally good poems on french fries

- Producing long poems, and maintaining coherency throughout them

Challenges

- Any user-supplied topic
  Like: burritos, Turing test, unicorn, Donald Trump

- Any number of distinct poems on that topic
  One million equally good poems on french fries

- Producing long poems, and maintaining coherency throughout them

- Small training data for each poem format

Bipolar Disorder

Existence enters your entire nation.
A twisted mind reveals becoming manic,
An endless modern ending medication,
Another rotten soul becomes dynamic.

Or under pressure on genetic tests.
Surrounded by controlling my depression,
And only human torture never rests,
Or maybe you expect an easy lesson.

Or something from the cancer heart disease,
And I consider you a friend of mine.
Without a little sign of judgement please,
Deliver me across the borderline.

An altered state of manic episodes,
A journey through the long and winding roads.

A Glance

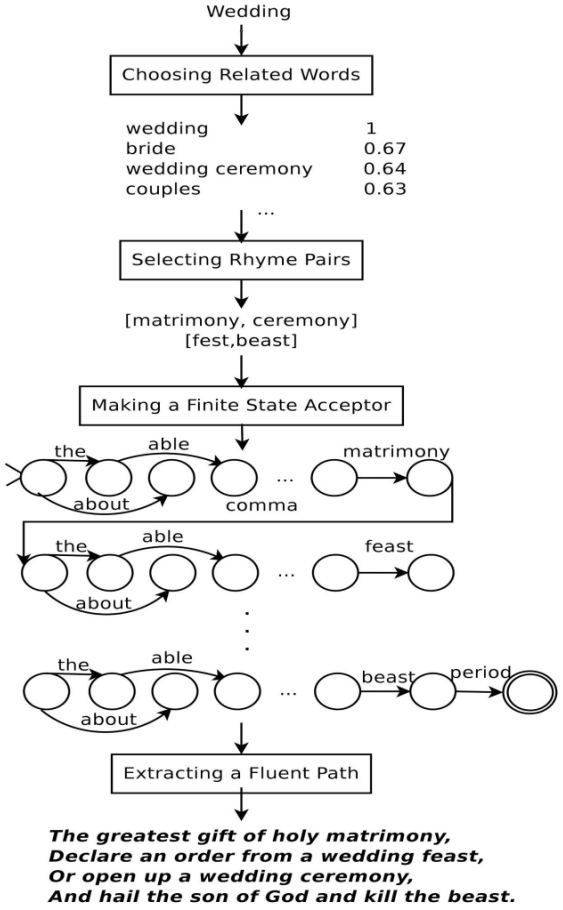
Wedding
Choosing Related Words
wedding 1
bride 0.67
wedding ceremony 0.64
couples 0.63
...
Selecting Rhyme Pairs
[matrimony, ceremony]
[fest,beast]
Making a Finite State Acceptor
the
able
about
comma
matrimony
feast
beast
period
Extracting a Fluent Path
The greatest gift of holy matrimony,
Declare an order from a wedding feast,
Or open up a wedding ceremony,
And hail the son of God and kill the beast.

Steps

➔Choosing related words

➔Selecting rhyme pairs

➔Making a finite state acceptor

➔Extracting a fluent path

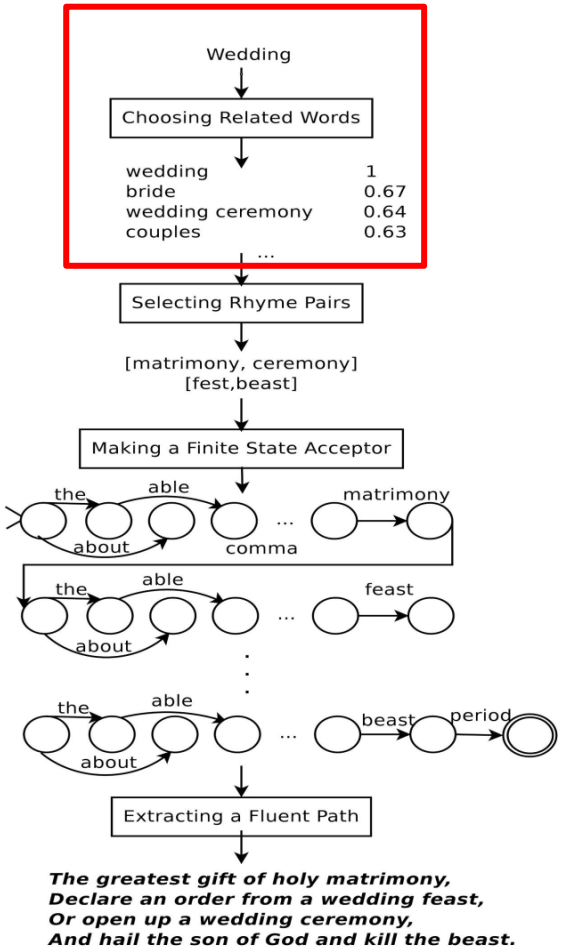

Steps

➔Choosing related words

➔Selecting rhyme pairs

➔Making a finite state acceptor

➔Extracting a fluent path

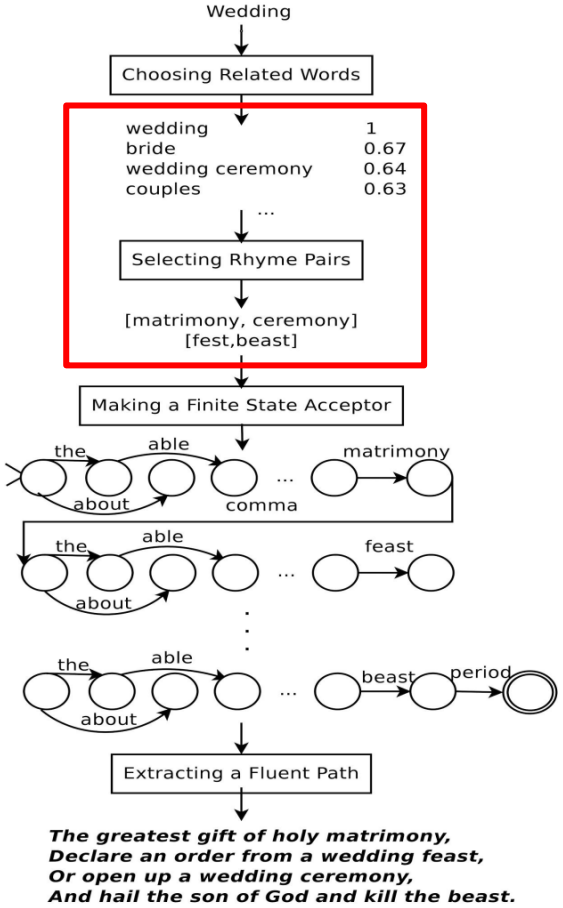

Steps

➔Choosing related words

➔Selecting rhyme pairs

➔Making a finite state acceptor

➔Extracting a fluent path

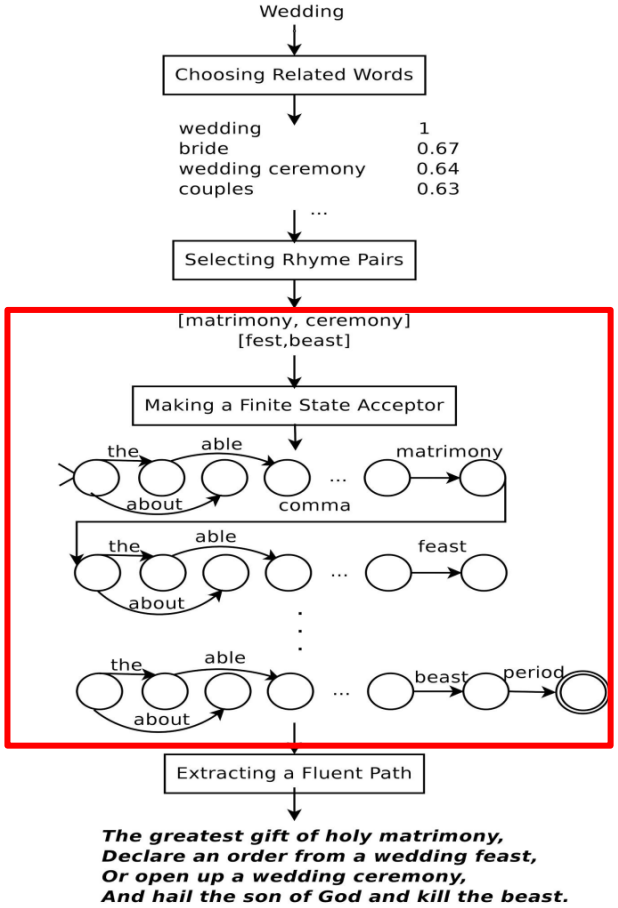

Steps

➔Choosing related words

➔Selecting rhyme pairs

➔Making a finite state acceptor

➔Extracting a fluent path

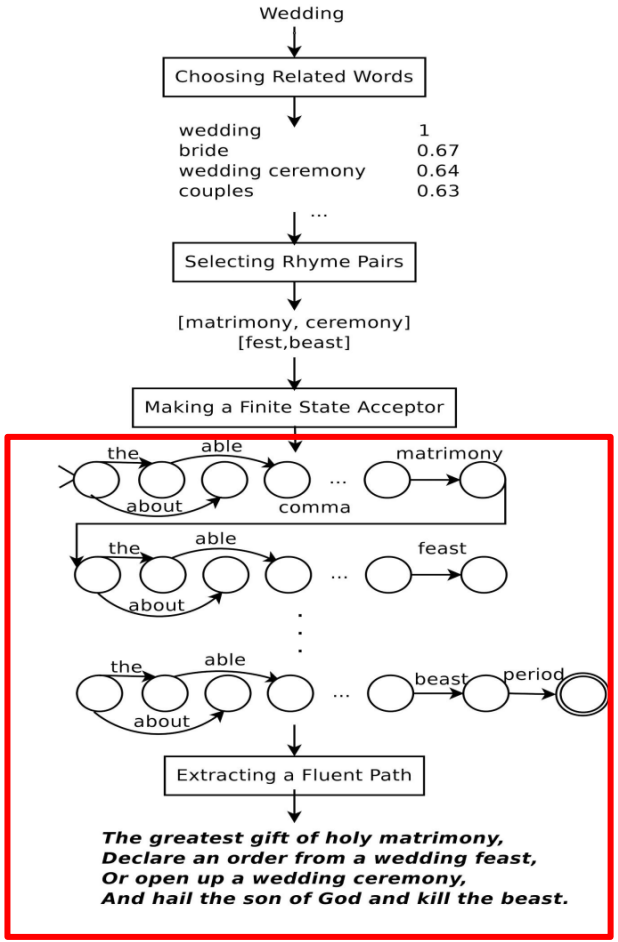

Choosing Related Words

Word2Vec:

Window size	Related words for *wedding*
8	wedding ceremony, marriage, ceremony, wedding celebration

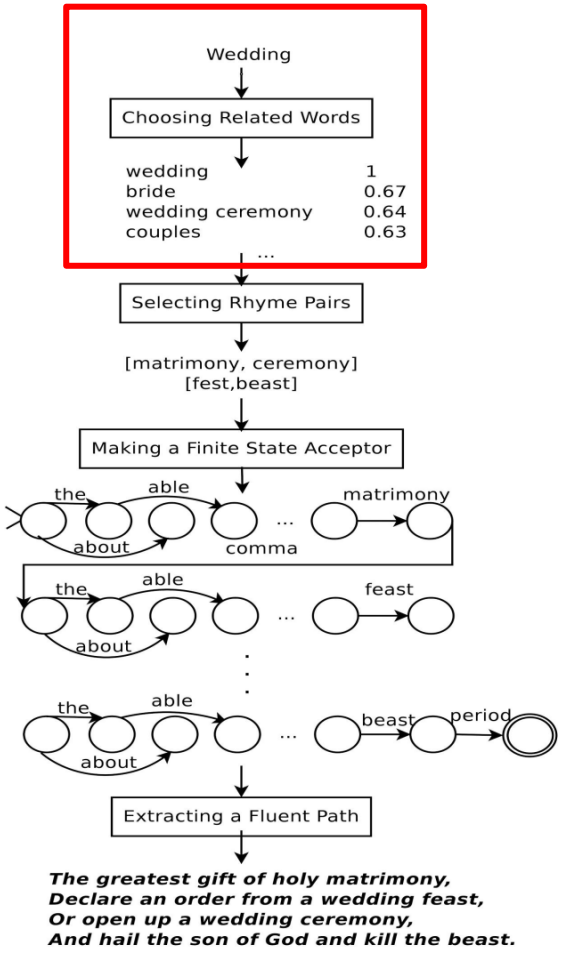

Choosing Related Words

Word2Vec:

Window size	Related words for *wedding*
8	wedding ceremony, marriage, ceremony, wedding celebration
40	bride, wedding ceremony, couple, weddings, marriage

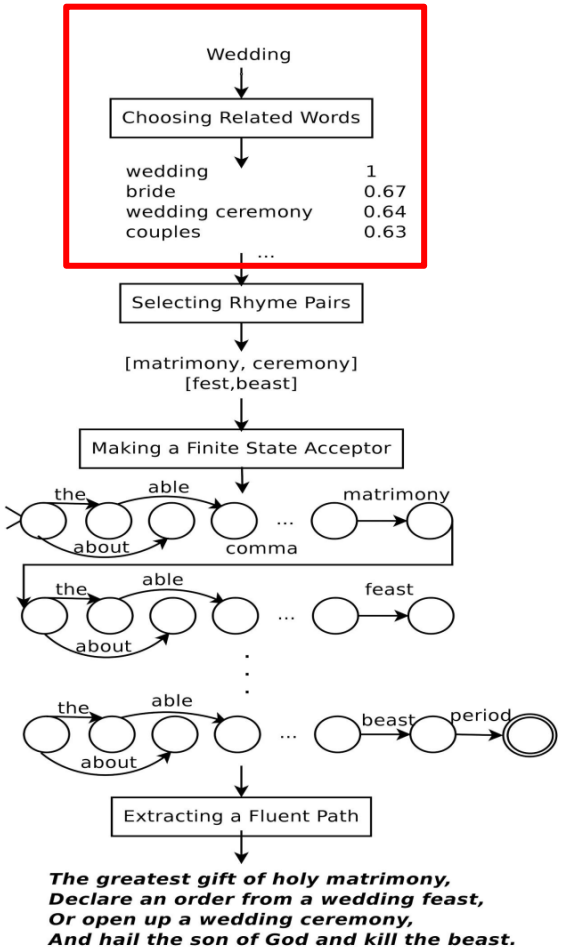

Choosing Related Words

Word2Vec:

Window size	Related words for *wedding*
8	wedding ceremony, marriage, ceremony, wedding celebration
40	bride, wedding ceremony, couple, weddings, marriage

Similarity score: Cosine similarity

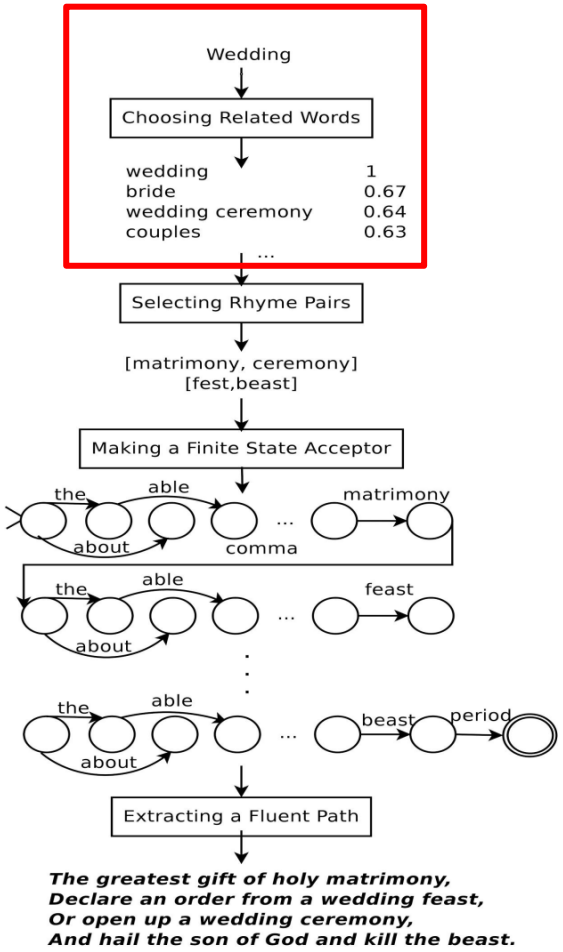

Selecting Rhyme Pairs

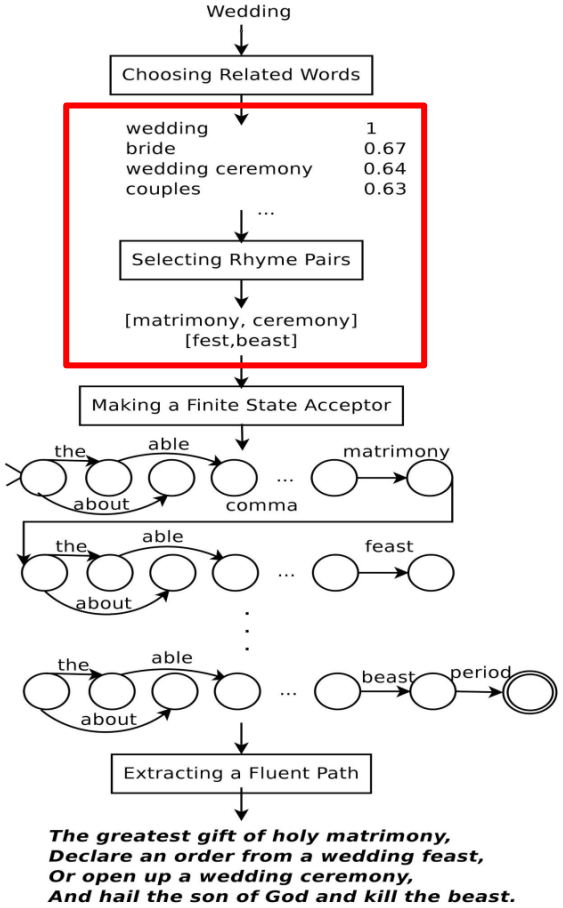

Selecting Rhyme Pairs

Check if two words rhyme together.

Exact rhyme:

Attend/friend, feast/beast, child/smiled

Slant rhyme

Viking/fighting, snoopy/spooky, comic/ironic

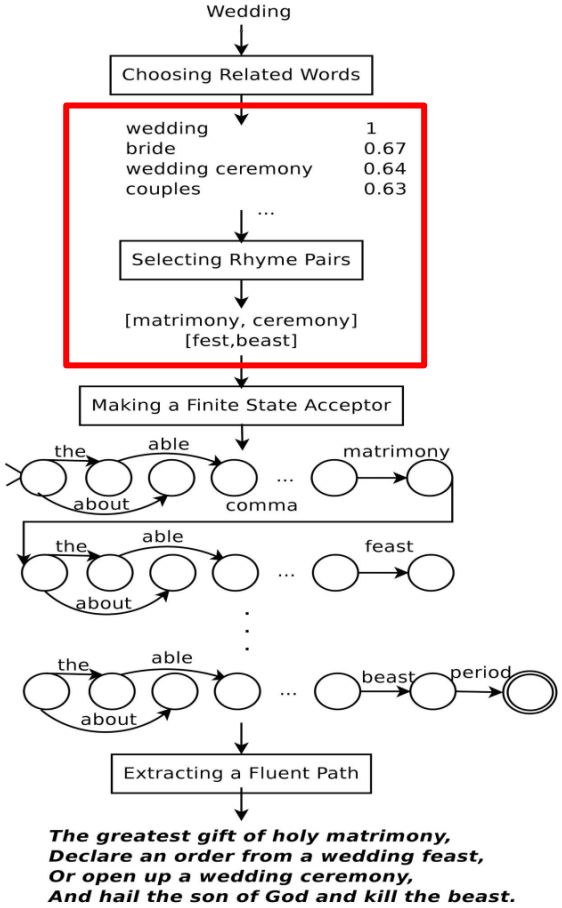

Selecting Rhyme Pairs

We select all rhyme pair candidates from the list of related words.

Topic: wedding

Candidate pairs: [dressed, guest] [celebration, invitation] [decorate, congratulate] [celebration, occasion] …

We choose rhyme word pairs randomly with probability proportional to the similarity of the words to the topic.

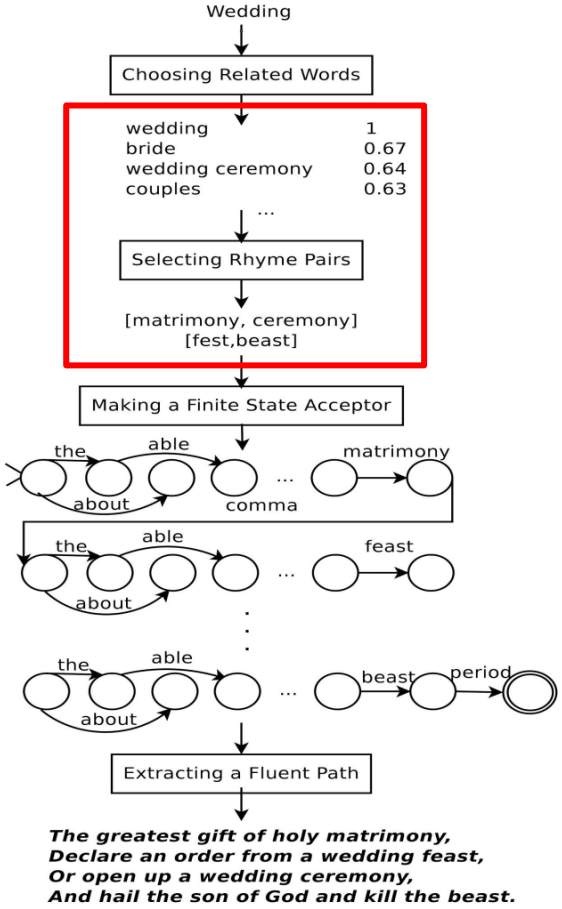

Making a Finite State Acceptor

Given rhyme words, we model all possible word sequences that use them and obey the poem format.

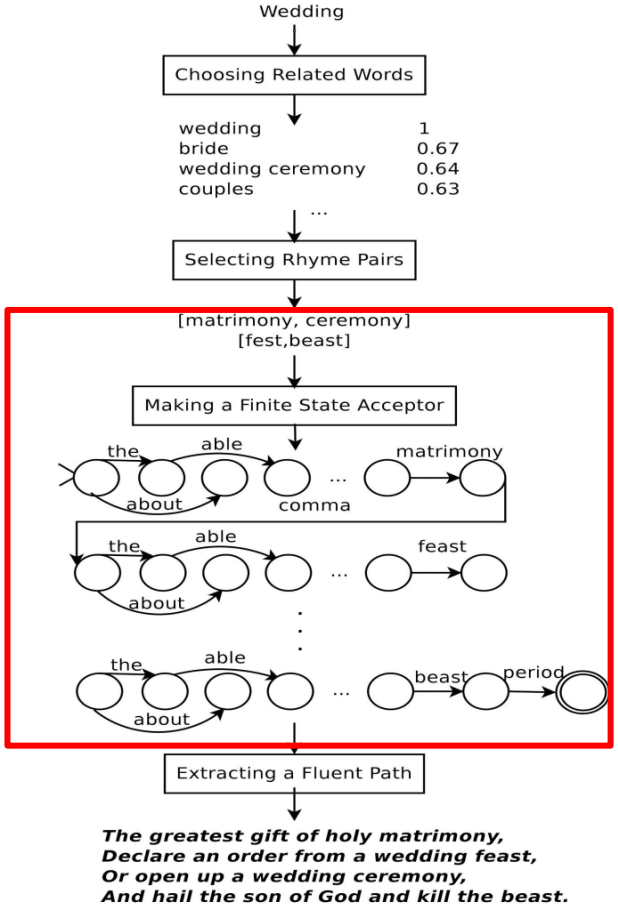

Format

Let’s assume we are generating 4-line stanzas.

Each line: A sequence of ten syllables alternating between unstressed and stressed.

*0:unstressed 1:stressed

Attending on his golden pilgramage

010 1 0 10 101

The greatest gift of holy matrimony,
Declare an order from a wedding feast,
Or open up a wedding ceremony,
And hail the son of God and kill the beast.

Format

Let’s assume we are generating 4-line stanzas.

Each line: A sequence of ten syllables alternating between unstressed and stressed.

*0:unstressed 1:stressed

Attending on his golden pilgramage

010 1 0 10 101

The greatest gift of holy matrimony

0 10 1 0 10 1010

The greatest gift of holy matrimony,
Declare an order from a wedding feast,
Or open up a wedding ceremony,
And hail the son of God and kill the beast.

Making a Finite State Acceptor

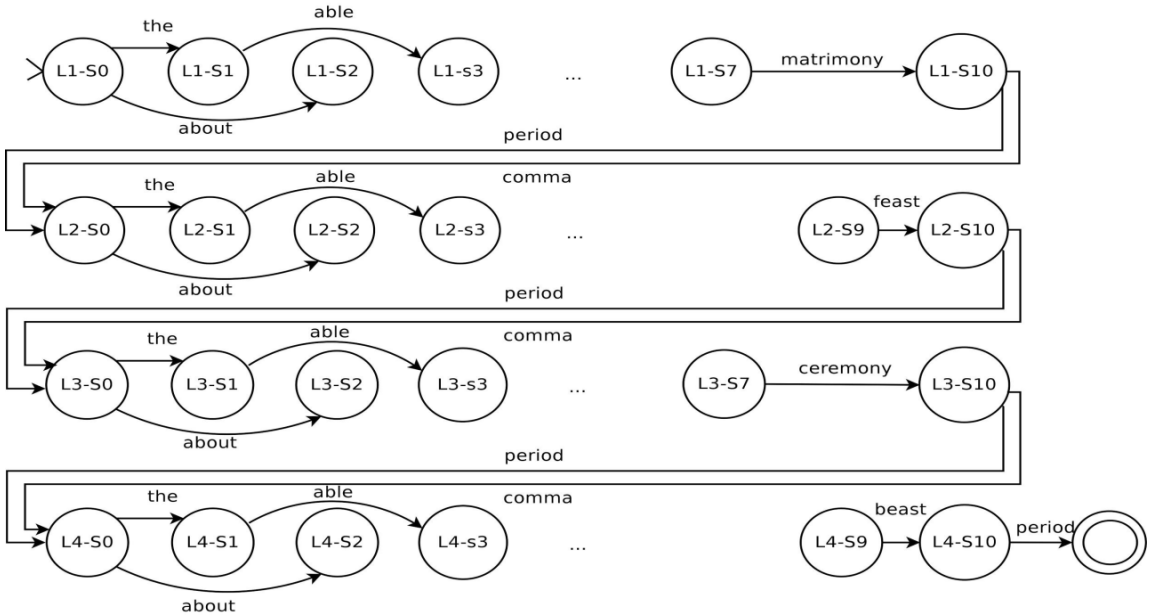

Making a Finite State Acceptor

FSA state L1-S3 signifies “I am in line 1, and I have seen 3 syllables so far”.

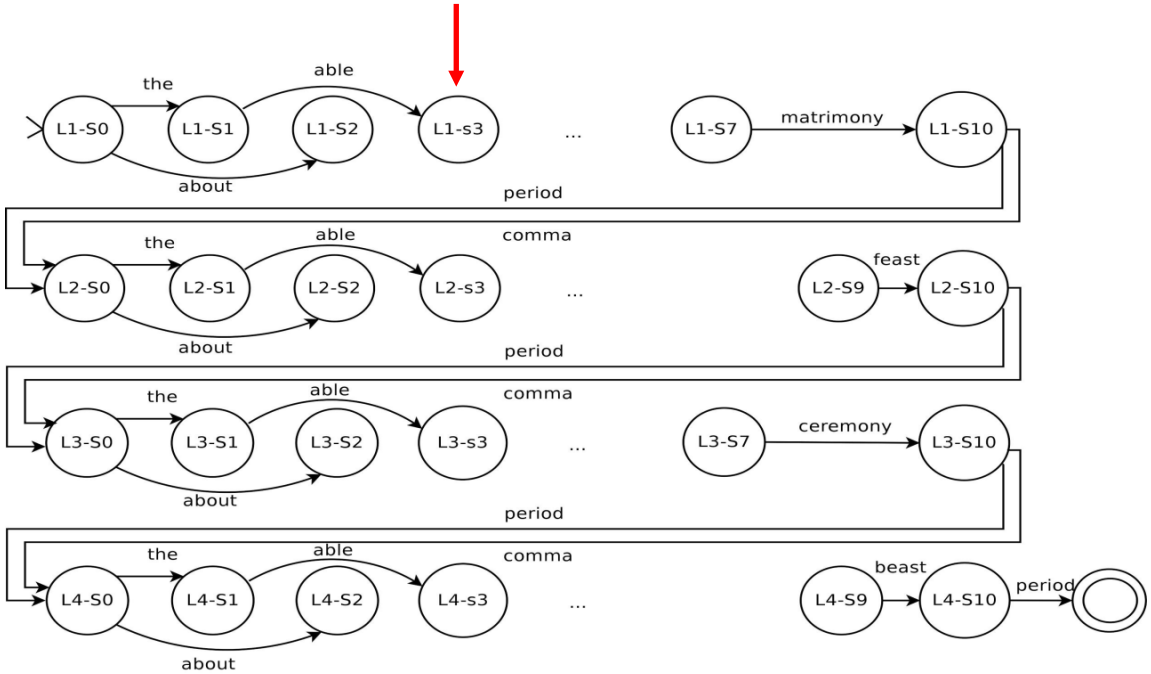

Making a Finite State Acceptor

FSA state L1-S3 signifies “I am in line 1, and I have seen 3 syllables so far”.

Format:

- 4 line
- 10 syllables
- 0101010101 pattern
- Fixed rhyme words

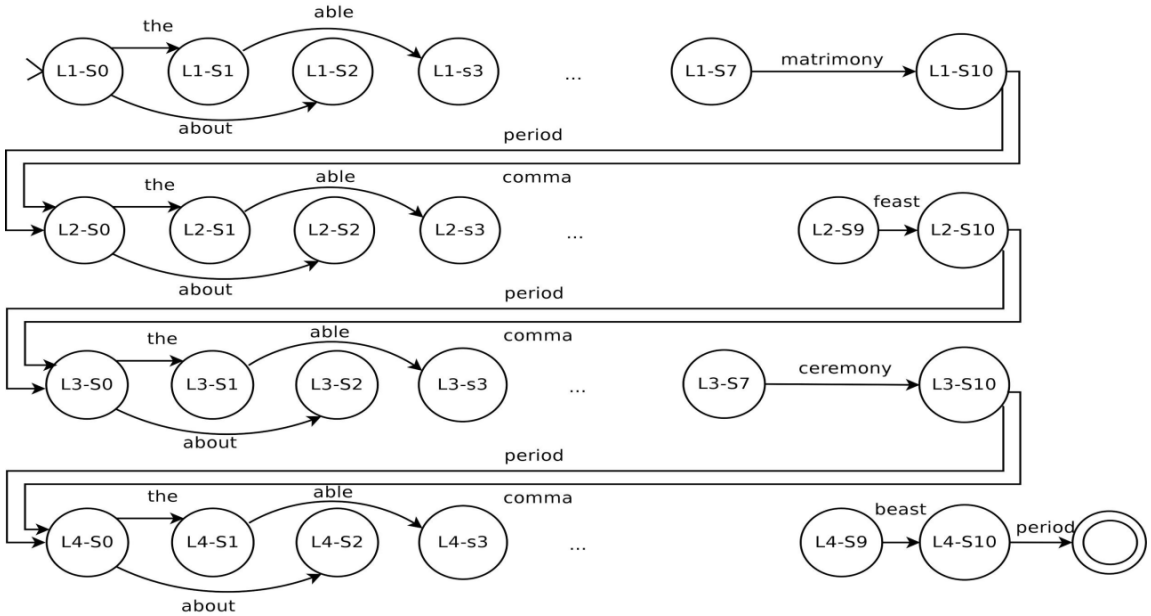

Making a Finite State Acceptor

This FSA contains 10^{67} 4-line possible poems.

Making a Finite State Acceptor

This FSA contains 10^{67} 4-line possible poems.

However most of them are meaningless:

> And roaring owners plaster matrimony.
> Or story Angie rail dimension feast,
> An hyper rented Walcott ceremony,
> Of Cochran grady fashion prison beast.

Extracting a Fluent Path

How to find a fluent path in this FSA?

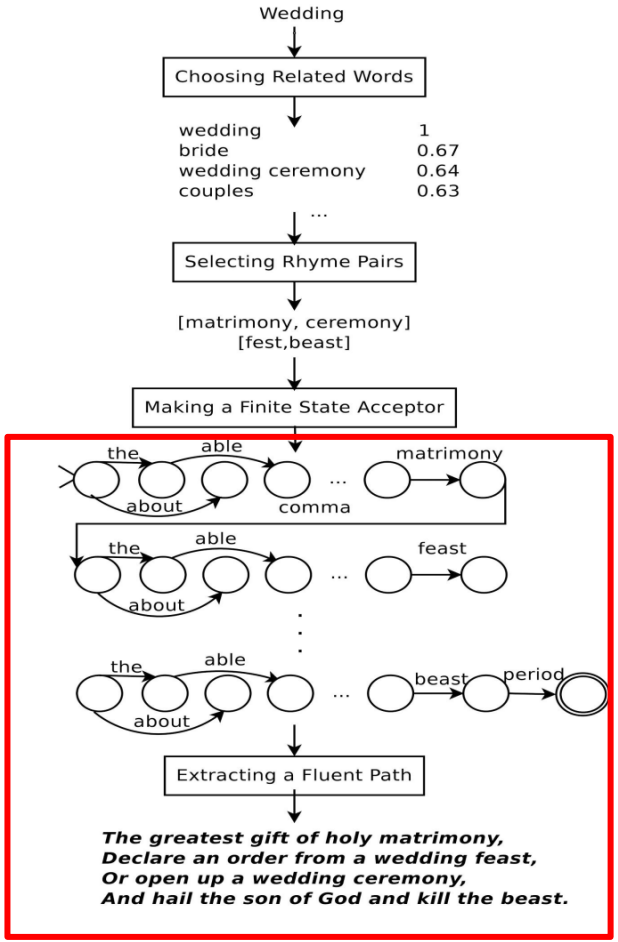

Extracting a Fluent Path

How to find a fluent path in this FSA?

Idea 1: N-gram LM as a large weighted FSA.

Intersect the two FSAs.

Find the path with the highest score.

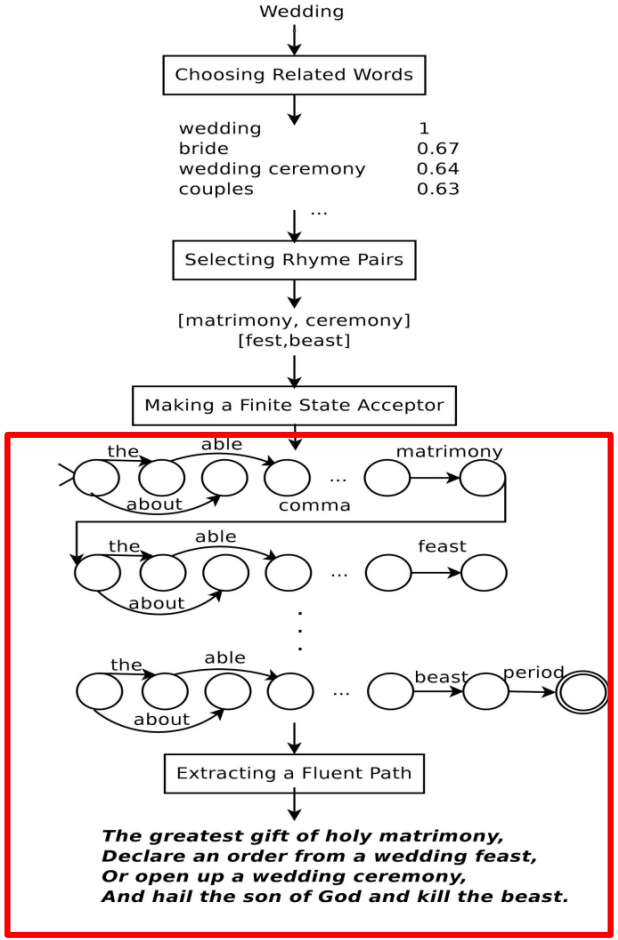

Extracting a Fluent Path

How to find a fluent path in this FSA?

Idea 2: RNN language model.

How to generate the right format?

How to intersect it with poem FSA?

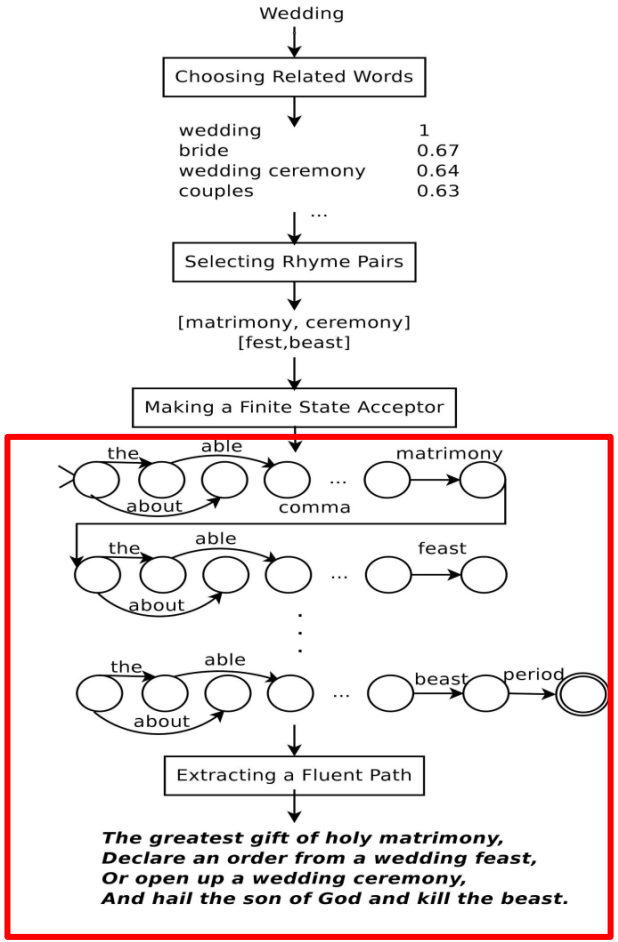

Extracting a Fluent Path Using RNN

- FSA accepts the word string with right format.
- RNN assigns a probability to each word string.

What we need:

The path through FSA with highest probability.

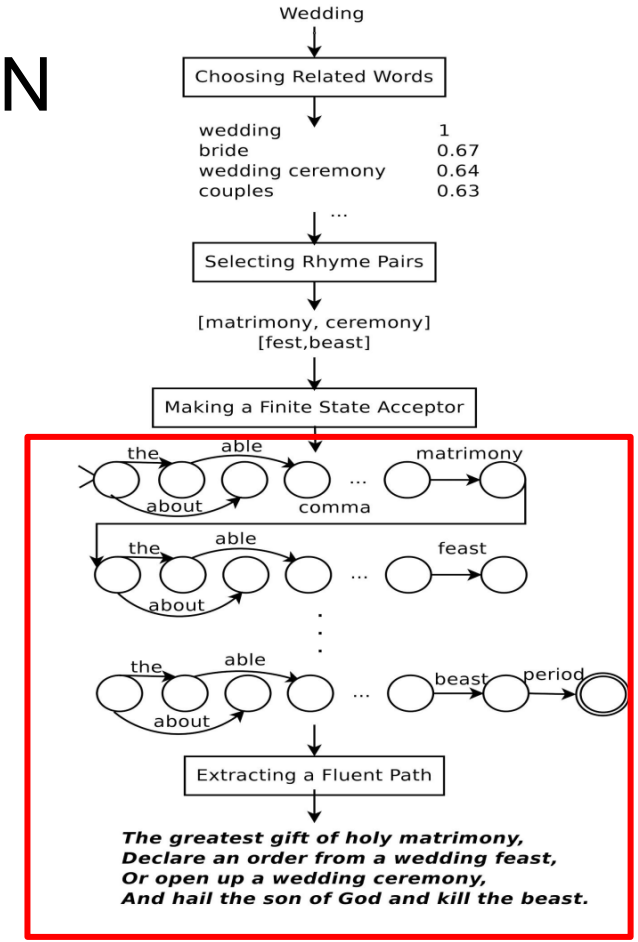

Extracting a Fluent Path Using RNN

- FSA accepts the word string with right format.
- RNN assigns a probability to each word string.

What we need:

The path through FSA with highest probability.

Intersection of RNN and FSA

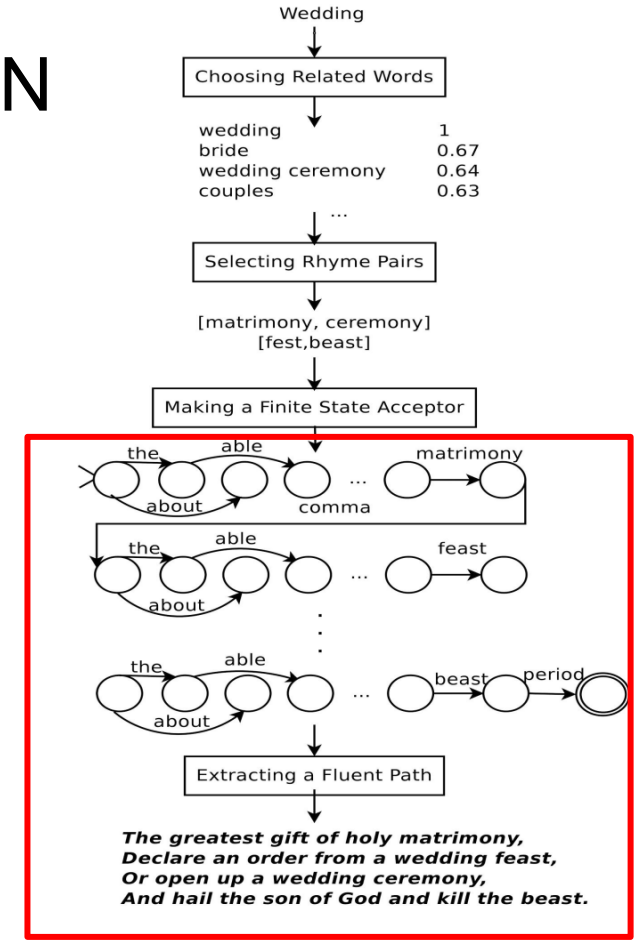

Intersecting RNN and FSA

RNN requires beam search to find a fluent path.

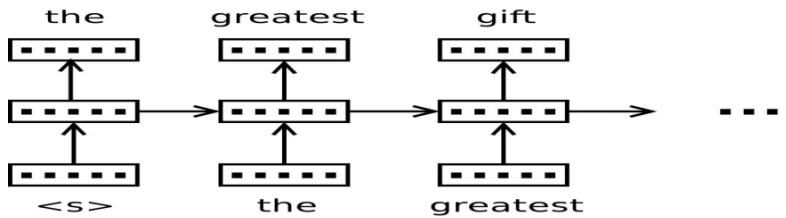

Intersecting RNN and FSA

RNN requires beam search to find a fluent path.

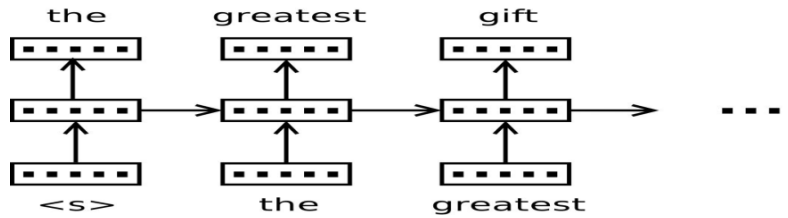

Intersecting RNN and FSA:

Same idea, employ a beam search that is further guided by the FSA.

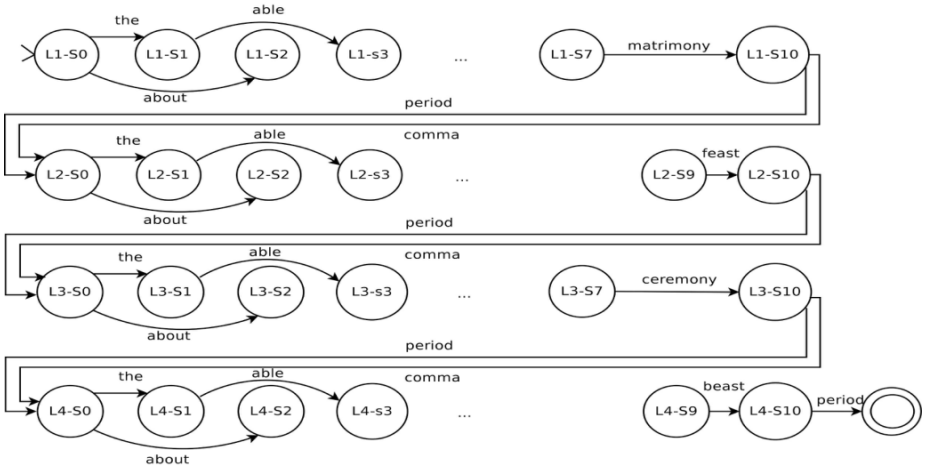

Vanilla RNN

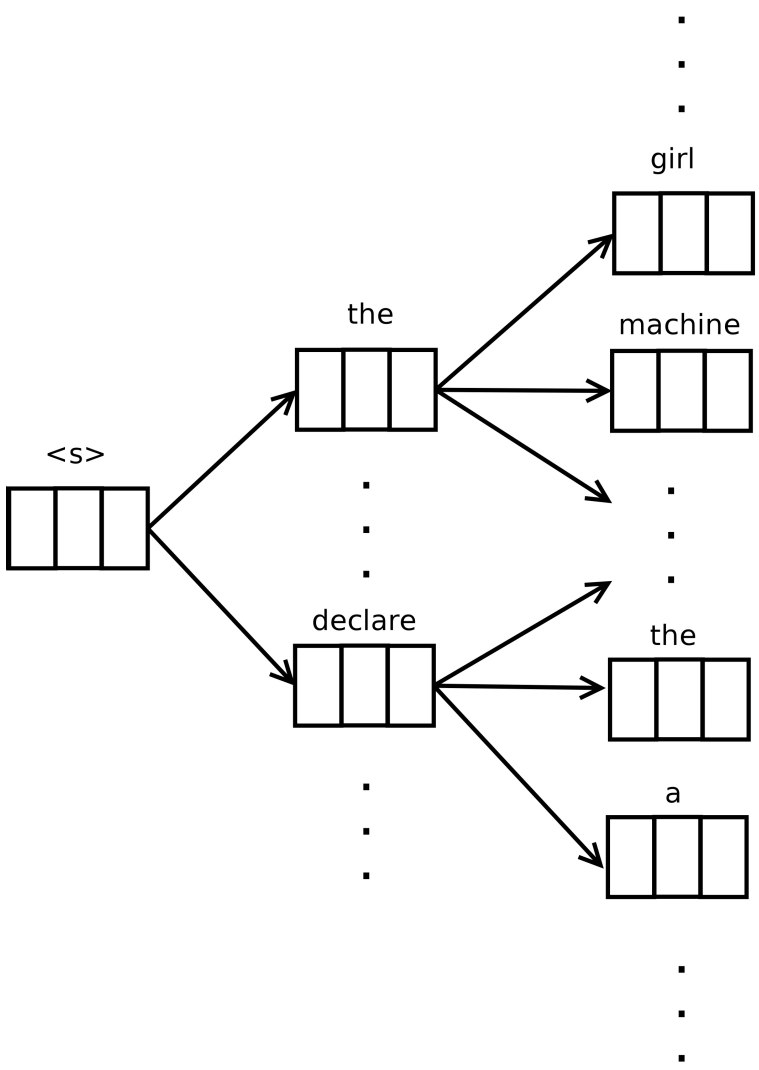

Vanilla RNN

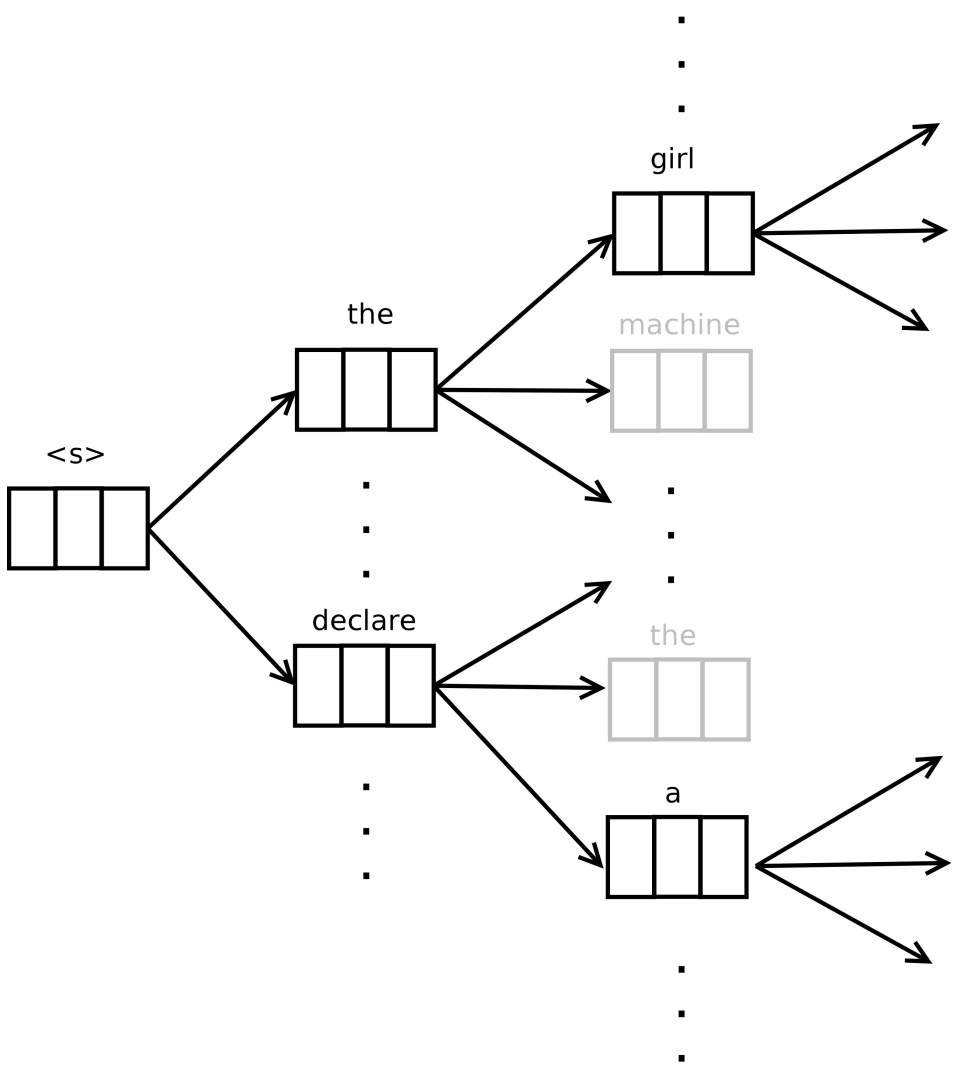

Intersection

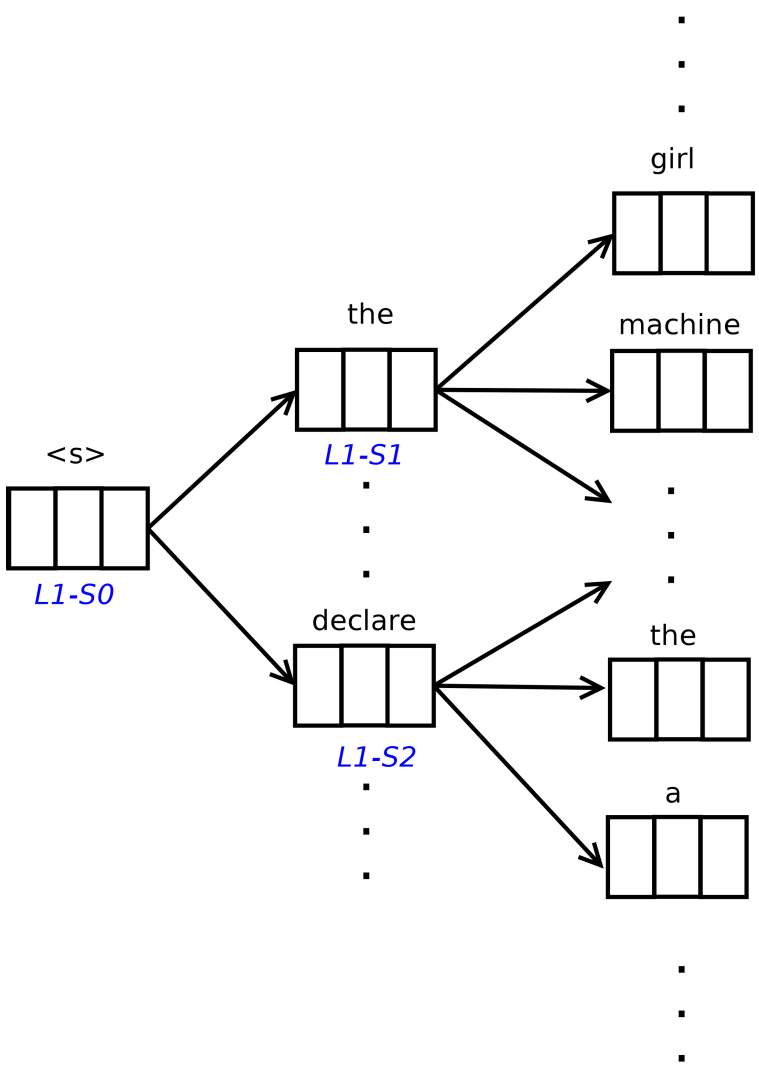

Intersection

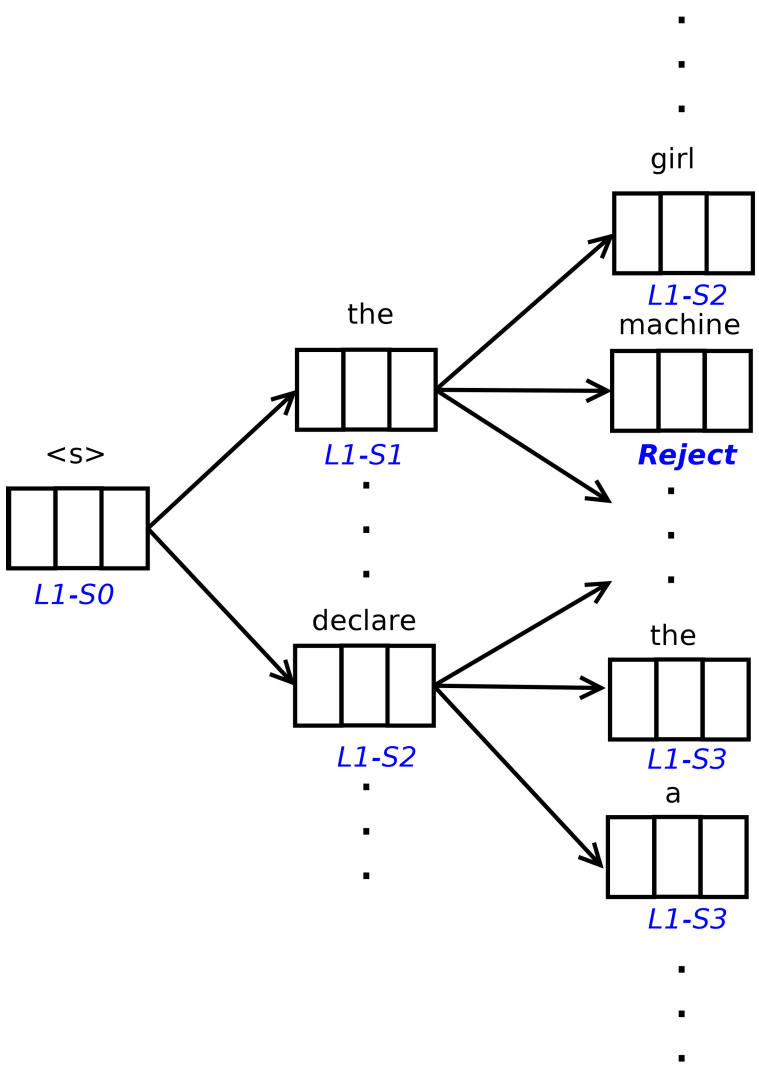

Intersection

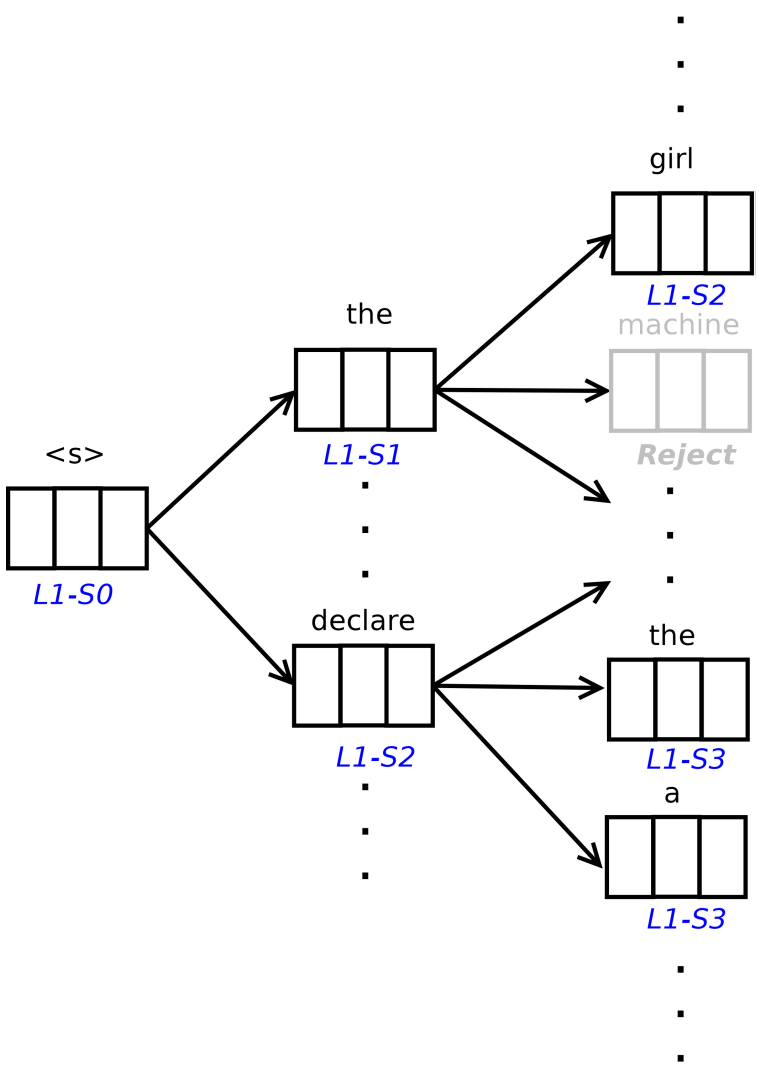

Extracting a Fluent Path

Now we can find a fluent path in the FSA.

In practice, We collect 94,882 English *songs* (32m word tokens) and train a 2-layer recurrent network with LSTM cells.
we choose beam size=50.

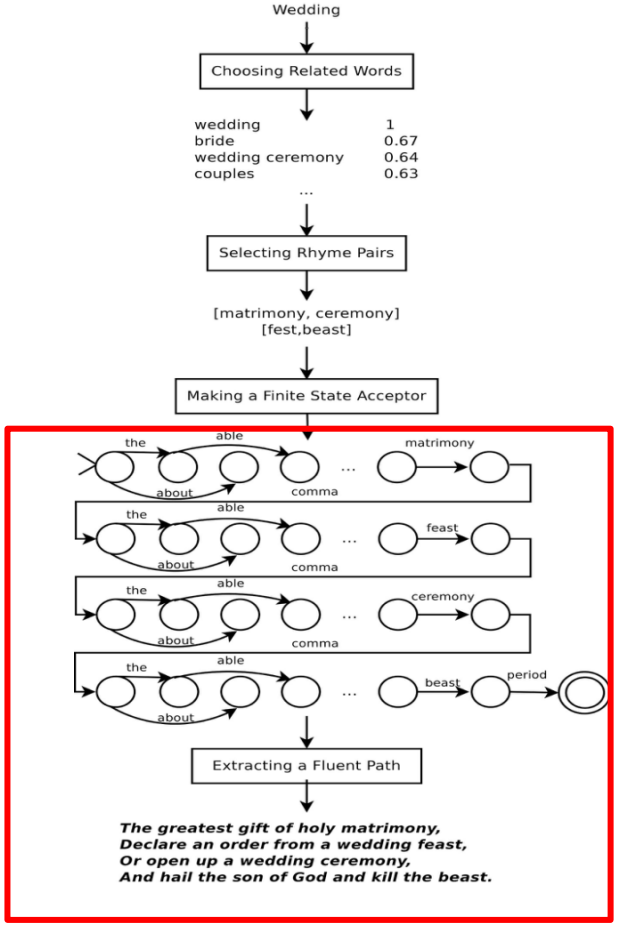

Beam search is not an optimal search

What if none of beam hypotheses can be meaningfully followed by the fixed rhyme word?

The greatest gift of holy matrimony,

Declare an order from a wedding feast,

Or open up a wedding ceremony,

And hail the son of God and kill the beast.

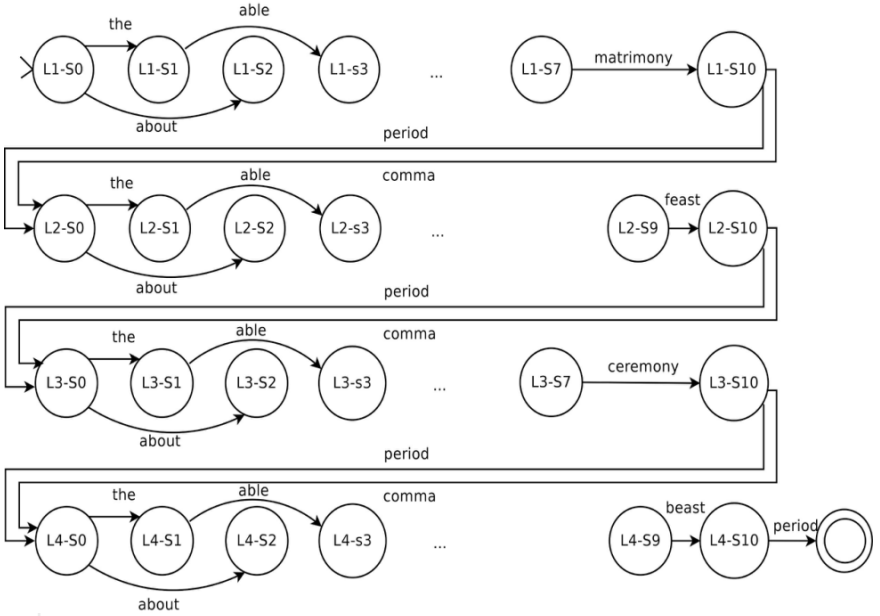

Beam search is not an optimal search

Solution: Generating poem in reverse

The greatest gift of holy matrimony,

Declare an order from a wedding feast,

Or open up a wedding ceremony,

And hail the son of God and kill the beast.

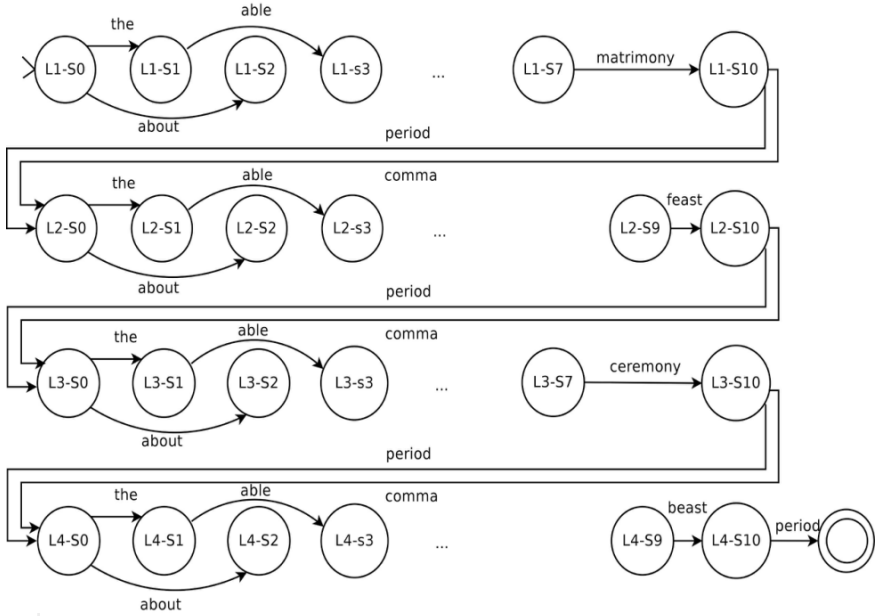

Beam search is not an optimal search

Solution: Generating poem in reverse

The greatest gift of holy matrimony,

Declare an order from a wedding feast,

Or open up a wedding ceremony,

And hail the son of God and kill the beast.

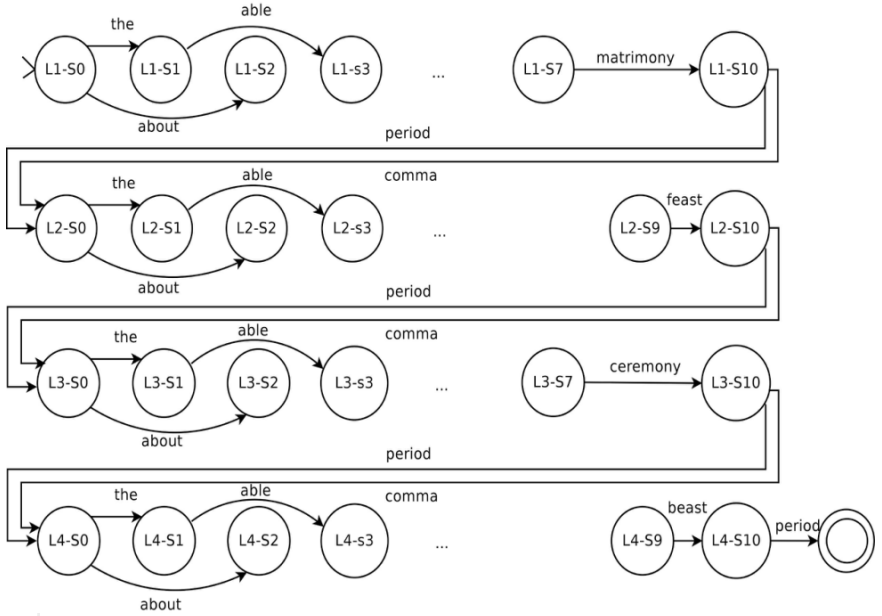

Beam search is not an optimal search

Solution: Generating poem in reverse

The greatest gift of holy matrimony,

Declare an order from a wedding feast,

Or open up a wedding ceremony,

And hail the son of God and kill the beast.

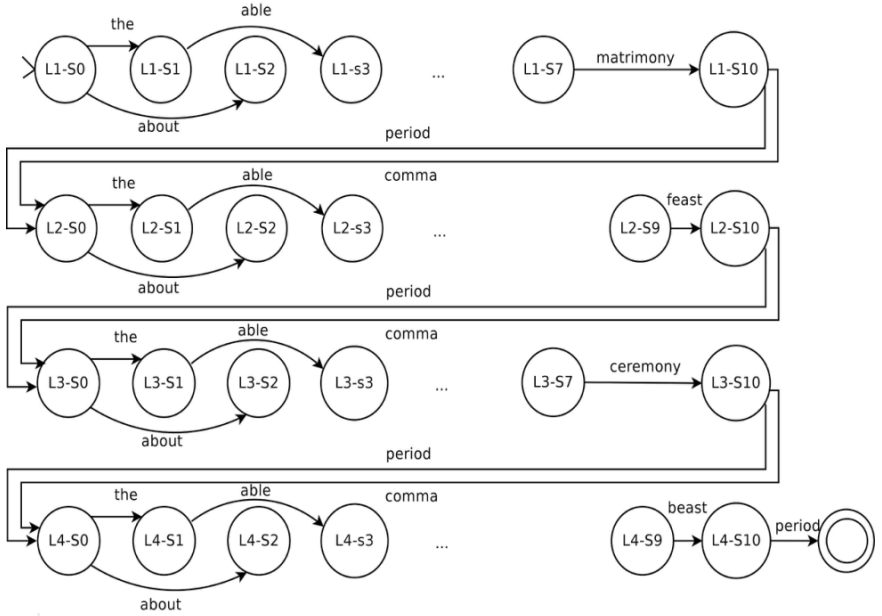

Beam search is not an optimal search

Solution: Generating poem in reverse

The greatest gift of holy matrimony,

Declare an order from a wedding feast,

Or open up a wedding ceremony,

And hail the son of God and kill the beast.

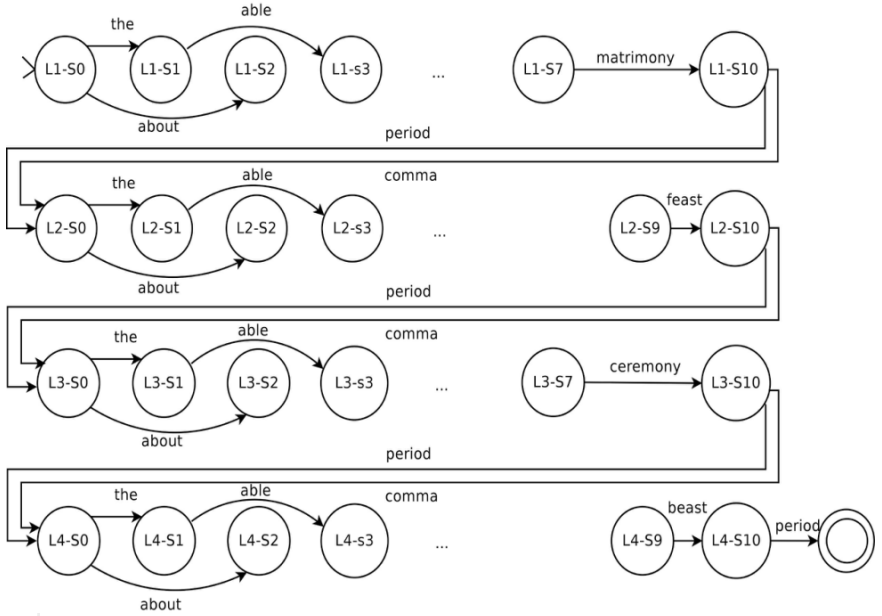

Beam search is not an optimal search

What if none of beam hypotheses meaningfully relate to all of the rhyme words?

The greatest gift of holy matrimony,

Declare an order from a wedding feast,

Or open up a wedding ceremony,

And hail the son of God and kill the beast.

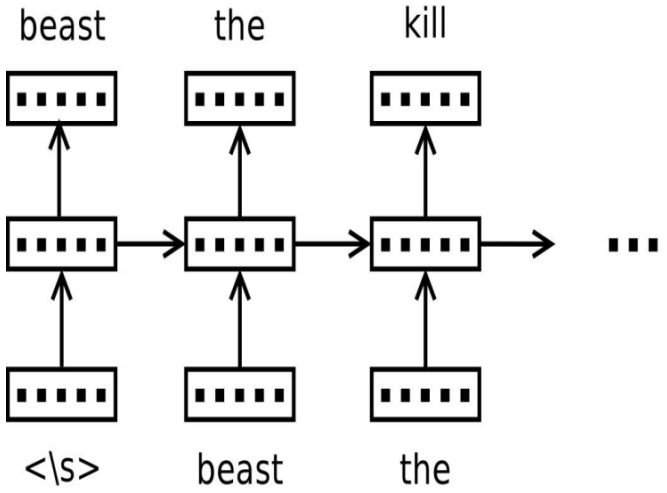

Beam search is not an optimal search

Solution: Providing all rhyme words to the Language Model

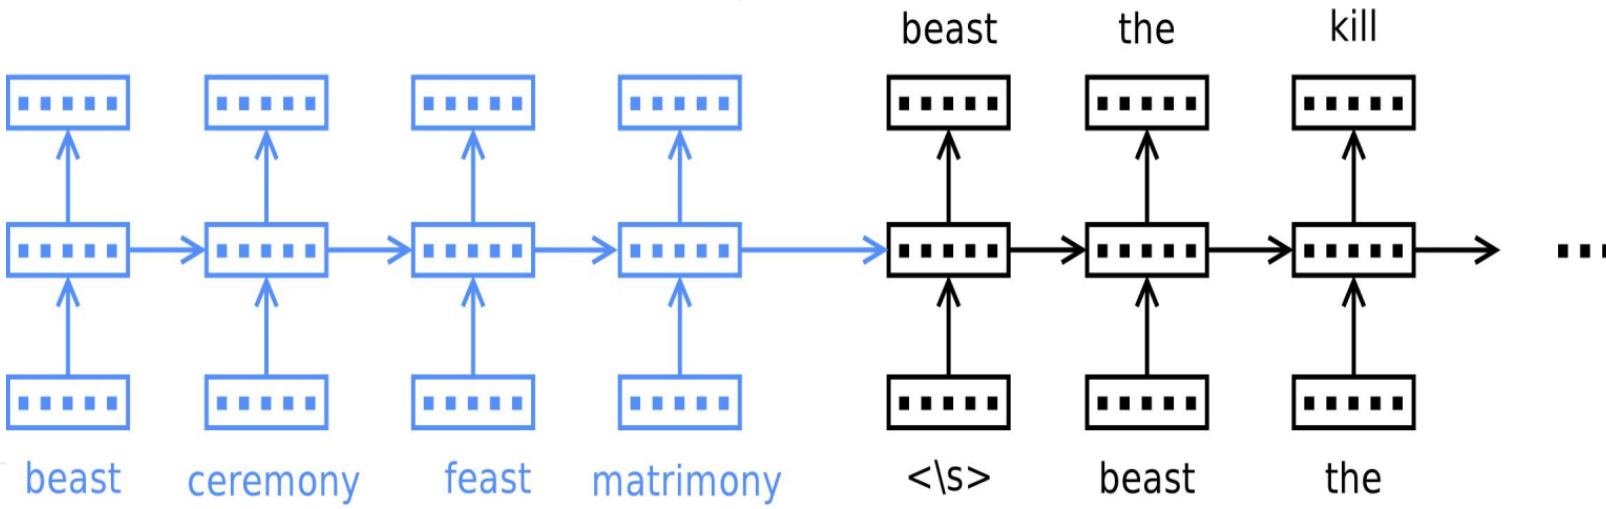

User Study

User preferences between generation model and translation model

Preference	Generation Model	Translation Model	Can not Decide
Stanzas	26%	**43%**	31%
Sonnets	21%	**57%**	22%

Sample poems

Love at First Sight

An early morning on a rainy night,
Relax and make the other people happy,
Or maybe get a little out of sight,
And wander down the streets of Cincinnati.

Sample poems

Noodles

The people wanna drink spaghetti alla,
And maybe eat a lot of other crackers,
Or sit around and talk about the salsa,
A little bit of nothing really matters.

Sample poems

Civil War

Creating new entire revolution,
An endless nation on eternal war,
United as a peaceful resolution,
Or not exist together any more.

Plagiarism or Creativity?

It is common for RNNs to repeat large sections of training data.

But our poetry system deals with lots of constraints.

- Rhymes
- Iambic rhythm
- Word-repeated penalty

Repeated 5-grams

	Repeated 5-gram from training data
Our system	3%
- word-repeated penalty	21%
- iambic rhythm	30%

Can It be Generalized?

Port to Spanish.

Classical Spanish soneta: 14 eleven-syllable lines under the rhyme scheme ABBA ABBA CDC DCD.

Mariposa

Quieres saber dónde está el escorpión,
Ni ayer ni antes vos sos corona dorada.
Ya os ves más tal cual tortuga pintada,
A él nos gusta andar con cola marrón.

Ella es quién son las alas de algún gorrión.
Si al fin podés ver tu imagen manchada,
O hoy vas bajo un cielo azul plateada,
Por qué estás tan lejos del aguijón.

No hay luz que al sol se enreda en tus palmera.
Ay por qué eres víbora venenosa,
Sin querer igual a un enredadera.

Y si aún sueñas con ser mariposa,
En vez de abrir los ojos y espera,
Sabes muy bien que el amor no es gran cosa.

Summary

We introduce a general method for combining finite state machinery with deep learning model.

We create Hafez: an end-to-end system for generating computer poetry from any user-supplied topic.

We show how to extend our system to different formats and languages.

Are We There yet?

Dartmouth

Topic word/phrase (e.g., "car keys")

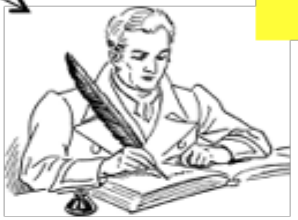

April 2016
- Very difficult, open-ended task
- Won $3000 first prize
- But automatic programs do not yet fool human judges!

Unbounded number of
14-line sonnets on the topic

Future Work

Modeling the point of the poem.

Importing more discourse information to the the system.

Using hierarchical deep learning models.

Thanks

ghazvini@isi.edu
http://isi.edu/~ghazvini/